

**DATA MIGRATION SYSTEM,  
METHOD AND PROGRAM PRODUCT**

**Forrest L. Wilson  
Jason J. Blackburn**

**BACKGROUND**

**Technical Field**

Our invention relates to migration of database information and files from one or more data storage systems to another data storage system, including backing up database information and files on one or more storage systems, and file migration to and from high density nonvolatile storage, with immediate, delayed, and scheduled backup.

**Description of Related Art**

Data intensive enterprises frequently operate with many different, disparate and dissimilar storage platforms from many different vendors. There is a frequent need to migrate data between these dissimilar storage platforms. As used herein “data migration” is the movement of data from one or more storage systems, that is, source storage systems, to target storage systems. Data migration may be motivated by upgrading storage systems, consolidating storage systems and data, replacing existing storage systems, and load balancing.

The data migration process is typically complex and labor-intensive in today’s business environment. This is because of the myriad of application servers, operating systems, file systems, volume management methodologies, physical devices, and networks.

Information Technology departments face ever more challenges in migrating data. These challenges include the downtime involved in data migration, the frequent need to add

data migration software to servers, the potential for data loss and corruption, the chance for error arising from heterogeneous environments, and simply the time involved.

Presently data is migrated between these different, disparate, and dissimilar platforms  
5 using host level mirroring with host based software. In the host based software approach with host level mirroring, the host software captures I/O operations before they leave the host and replicates them on the target storage device. But, installing host based software is itself inconvenient, requiring IPL's or reboots. Moreover, the host based software approach frequently requires additional software between the host operating system and  
10 the physical devices being migrated. Thus, host level mirroring with host based software is not a completely satisfying solution.

Another reason that host level mirroring is not altogether satisfactory is that many hosts do not supply the mirroring software, and therefore third party software solutions must be  
15 purchased. Typical third party solutions include Doubletake for Windows, HP Mirror, and Veritas for Sun. Data migration using third party software typically utilizes host level mirroring and is frequently labor intensive

Given these obstacles, a clear need exists for a data migration system, method, and  
20 program product that is non-intrusive to the underlying business process, and relatively easy to implement.

## SUMMARY

25 According to the method and system of our invention, it is possible to move data between disparate storage platforms, while avoiding host level mirroring. An embodiment of the present invention migrates computer data files as bit images of logical volumes between a

source data storage device and a target data storage device. These data storage devices may be disparate and dissimilar devices.

5 An embodiment of the present invention requests data from a source volume as a bit image of the volume and writes the data to a target volume as a bit image. The data is requested and, optionally, read as a bit image of a logical volume, cylinder by cylinder, track by track, and bit by bit. The data is moved on as an image of the logical source volume, on a cylinder by cylinder, track by track, and bit by bit basis. According to an embodiment of the present invention, this data is then written to a target volume on the  
10 target data storage device as a bit image of a logical volume, on a cylinder by cylinder, track by track, and bit by bit basis.

The data is migrated as logical volumes in accordance with a map file having source and target volume parameters. Generally, the logical volume comprises a physical volume.  
15

Using an embodiment of the present invention, a data migration hardware appliance, referred to herein as a data migration hardware engine, is connected between the existing storage devices, such as a host system and one or more target devices, data is transferred from source to target. The particular data transfer connection depends on the type of host  
20 server and storage device, (i.e., SCSI, SSA) open systems or mainframes.

The method, system, and program product described herein enable a user to access data from the source volume and move off of the source volume at substantially the same time.  
25

## **BRIEF DESCRIPTION OF THE DRAWINGS**

FIG. 1 is a flow chart of an embodiment of the invention for copying allocated tracks as a bit image from a logical source volume, and writing the bit image from the logical source volume to a logical target volume. The flow chart further shows that after migration is complete, I/O is switched from the source to the target.

5

FIG. 2 is a high level schematic overview of a system configured to carry out an embodiment of the present invention. The FIG. shows a source platform, a data migration engine, and a target platform. Also shown are data updates, and data migration pathways, with one pathway from the source platform to the data migration engine and then from the data migration engine, to the target platform.

10

FIGs 3A, 3B, and 3C, illustrate a data migration sequence where servers are initially utilizing a first data storage platform to store data. Figure 3B shows the installation of a data migration engine to migrate data from the initial or source storage system to a target storage system. Figure 3C shows the servers now using the target storage system after data migration.

15

FIG. 4 illustrates a data migration system for a mainframe computer system, with two mainframe computers, two data storage systems, a data migration engine, and fiber optic connector.

20

FIG. 5 illustrates an example of a signal-bearing medium carrying a program product to configure and control a system in accordance with an embodiment of the invention.

25

## **DETAILED DESCRIPTION**

As described herein, data is moved between disparate storage platforms, while avoiding host level mirroring. An embodiment of the present invention migrates computer data

SJO920030075US1

files as cylinder-by-cylinder, track-by-track, bit-by-bit images ("bit images") between a source data storage device and a target data storage device. These data storage devices may be typically disparate and dissimilar devices.

5

## OVERALL SEQUENCE OF OPERATION

For ease of explanation, but without any intended limitation, one aspect of the invention is described with reference to the flow chart illustrated in FIG. 1, and the system shown generally in FIG. 2.

10

The method and system of the invention migrates computer data files between a source data storage device and a target data storage device, where the computer data files are accessible to an end user from either data storage device. This is done without host level mirroring, requesting, and optionally reading, data from a source volume on the source data storage device, as a bit image of a logical volume; and writing the data to a target volume on the target data storage device, as a bit image of a logical volume.

15

FIG. 1 is an illustration of a flow chart of one embodiment of the invention. The illustrated embodiment starts data migration by starting a monitor task on systems, 17, having access to storage, block 1. Block 2 shows beginning the migration of the data in the source volume on the source storage, 11, to target storage, 15. The embodiment of the invention copies allocated tracks as bit images from the logical source volume, as shown in block 3. This embodiment writes the allocated tracks as a bit image from a logical source volume to a logical target volume, as shown in block 4. In one embodiment, illustrated in block 5 of the flow chart of FIG. 1, updates to the source are detected and the updated tracks containing updates to the source are recopied as a bit image from the logical source to the target. After migration is complete, I/O is switched from the source to the target.

20

25

This embodiment of the invention facilitates requesting, and optionally reading, data from a source volume and writing the data to a target volume. The data is read as a bit image of a logical volume, cylinder by cylinder, track by track, and bit by bit. The data is then moved on as an image of the logical source volume, on a cylinder by cylinder, track by track, and bit by bit basis, and written to a target volume on the target data storage device as a bit image of a logical volume, on a cylinder by cylinder, track by track, and bit by bit basis. The data is migrated as logical volumes in accordance with a map file having source and target volume parameters. Generally, the logical volume comprises a physical volume.

During the migration process, each source volume is mapped to a target LUN, (as used herein an "LUN" is a logical unit number on a SCSI bus that allows differentiation between and addressing to up to eight logical units). creating a synchronous mirror between each pair. All write commands go to both members, while completing a copy to the target volumes. As with any mirroring scenario, updates during the synchronization process are written to both sides of the mirror, allowing migration to continue with concurrent access from the applications. Once the data migration is complete, and the synchronous mirrors are in sync, a hard busy is placed on the source volumes and the SCSI ID from the source volumes are moved to the target volumes. This process allows the data migration with the computer data files being accessible to an end user during the data migration.

In this way a user accesses data from the source volume and moves off of the source volume at substantially the same time.

#### HARDWARE COMPONENTS AND INTERCONNECTIONS

The hardware embodiment of the invention includes a data migration engine, 13, connected between the existing or source data storage devices, the host system (server) and the new or target data storage devices. The type of connection depends on the type of host server and storage device, (i.e., FC, FCCAL, SCSI, SSA, ESCON, FICON), that is,  
5 open standards systems or mainframes.

FIG. 2 is a high level view of a simplified system, while FIG 3 is a schematic view of an open standards system, and FIG. 4 is a schematic view of a mainframe system. In all of these systems a data migration engine (also referred to herein as a data migration engine,  
10 a data migration appliance, and a data migration hardware appliance).

The embodiment of the present invention, illustrated in FIG. 2, requests, and optionally reads, data from a source volume on a source data storage device, 11. The data is obtained as a bit image of a logical volume, cylinder by cylinder, track by track, and bit  
15 by bit. The data is moved to a data migration engine, 13, such as an Innovation Data Processing, Inc. FDRPAS (FDR Plug and Swap) tool in a mainframe system or a Vicom System DME Data Migration Engine for open systems. The image is migrated as an image of the logical volume, on a cylinder by cylinder, track by track, and bit by bit basis, and from the data migration engine, 13, to a data target, 15, also, as an image of the  
20 logical volume on the source platform, 11, on a cylinder by cylinder, track by track, and bit by bit basis.

According to an embodiment of the present invention, this data is written to a target volume on the target data storage device, 15, as a bit image of a logical volume, on a  
25 cylinder by cylinder, track by track, and bit by bit basis.

The data is migrated as individual logical volumes in accordance with a map file having source and target volume parameters. Generally, the logical volume comprises a physical volume.

- 5 Also shown are data updates, 21a, 21b, and 21c, and data migration pathways, with one pathway, 23a, from the source platform, 11, to the data migration engine, 13, and then from the data migration engine, 13, to the target platform, 15, along migration path 23b.

- 10 Using an embodiment of the present invention, it is possible to receive updates during migration and maintain two synchronous copies of the data. Once the required logical volumes are synchronized, primary access to the data can be placed on the target volumes. This involves receiving updates during migration, and then placing a “busy” (e.g., a “hard busy”) condition on the source volume after data migration, and setting a SCSI ID to identify the target volume for access. This process is repeated, for example,  
15 on a logical volume by logical volume basis. In this way a user accesses data from the source volume and moves off of the source volume at substantially the same time.

- As used in the method and system of our invention, the data migration engine, 13, reads every track, every cylinder, and every bit, from the source platform, 11, and writes this  
20 information to the target platform, 15. The data being migrated is a bit copy, that is, a bit-by-bit, track-by-track, cylinder-by-cylinder image, of the original, source volume on the source platform, 11. Zeros and ones are moved from a source platform, 11, to a target platform, 15. For this reason, an embodiment of the present invention can migrate image databases, such as Oracle, IBM DB2, and Microsoft SQL7, without concern for data  
25 corruption or missing data.

## OPEN STANDARDS SYSTEMS



FIGs 3a, 3b, and 3C illustrate open standards systems in a data migration sequence. The FIGs illustrates open standards servers 17a and 17b. These servers may be Windows, Unix, Linux, or MacOS servers. The servers, 17a and 17b, are connected to optical connectors, 14, exemplified by Fibre Channel connectors. Fibre Channel is an optical  
5 channel for Gbps level data transmission between servers and, for example, shared storage devices, and for interconnecting storage devices and drivers.

In FIG 3a, the servers, 17a and 17b are connected directly to the data storage system 11. In FIG. 3b data migration is started from a source data storage system, 11 to a target data  
10 system, through a data migration system, 13, as will be discussed below. After data migration is completed, the servers, 17a and 17b, are disconnected from the source data storage system, and the data migration engine, 13, is removed from the system; the servers are connected through the Fibre Channel connector, 14, to the data storage system, 15.

15 For open systems, the data migration engine, 13, may be, strictly by way of illustration and not limitation, a Vicom Systems DME data migration engine. The Vicom DME data migration engine is an in-band appliance that provides block level migration between a source and a target storage device. The Vicom DME data migration engine maps each  
20 source volume to a target LUN. This tool transparently moves active volumes from old devices to new devices. Typically, the data migration engine is used in an environment that supports dynamic plug in and nondisruptive activation of new disk hardware.

For open systems, the Vicom DME data migration engine serially or concurrently swaps  
25 source devices SCSI ID's to the target devices while maintaining the individual source platform product data. This allows active open systems to continue to utilize existing drivers and, even, vendor specific software, while communicating with the target devices as if they are the original source devices.

The data migration engine, 13, is connected between the source platform, 11, and the target platform, 15. This connection can be a point to point configuration, an arbitrated loop configuration, or a SCSI configuration. The data migration engine, 13, is a pass-through between the source platform, 11, and the target platform, 15, and allows all updates to go to both the source platform, 11, and the target platform, 15. Both the source platform, 11, and the target platform, 15, receive updates through the data migration engine, 13.

10 The data migration engine, configured and controlled to read source data as an image of the logical volume, on a cylinder by cylinder, track by track, and bit by bit basis, and transfer it to the target platform, 15, starts reading source data from, for example, cylinder "0" and track "0", copying the data from the beginning of the source volume to the end of the source volume on a physical volume level.

15 After the migration or copy has been successfully completed, the data migration engine, 13, puts a "busy" (e.g., a "hard busy") condition on the source volume and switches the SCSI ID so that the target volume is the only volume online for the customer to use. This process is continued until all of the source volumes have been copied and switched to the target platform, 15.

For open systems utilizing, for example, it is not always necessary for the data migration engine, 13, to put a "busy" (e.g., a "hard busy") condition on the source volume and switch the SCSI ID so that the target volume is the only volume online for the customer to use. This is because, in the case of data migration engines for use with open systems, the tool, e.g., a Vicom DME, has the ability to maintain the source and target devices until such time as the host I/O for that particular open system can be stopped allowing for

the source and target to be exact duplicates of each other and maintain two copies of the data.

## MAINFRAME SYSTEMS

5

The data migration engine, 13, may be, strictly by way of illustration and not limitation, an Innovation Data Processing, Inc. FDRPAS (FDR Plug and Swap) data migration engine for mainframe systems. This engine transparently moves active volumes from old devices to new devices. Typically, the data migration engine is used in an environment that supports dynamic plug in and nondisruptive activation of new disk hardware. By way of illustration for mainframe, IBM z/OS system services in conjunction with FDRPAS, facilitates serially or concurrently swapping multiple shared disk volumes, in multiple systems.

15 Data migration can be carried while the end user is accessing old storage and also moving of the old storage at the same time. The data migration can be fully automated from a map file which has the source volume and the target volume in the map file. The map file is based on machine identifiers, as serial numbers.

20 The Innovation Data Processing, Inc. FDRPAS (FDR Plug and Swap) data migration engine for mainframe systems begins migration by a job request or console command. The command requests that an online disk storage device, the source device, 11, be migrated to an off line disk storage device, the target device, 15. The migration can be initiated on any server, 17c, in a complex, and other servers, 17d, will join in. The migration tool, 13, will assure that the target device, 15, is off line to all servers, 17c, and 25 17d, and can not be accidentally overlaid during data migration.

A fiber optic connector, 14, such as an IBM dynamically modifiable fiber optic interconnect, connects the

For all requested volumes, the Innovation Data Processing, Inc. FDRPAS (FDR Plug and Swap) data migration engine, 13, will copy all allocated tracks on the source volume, 11, to the target device, 15, while simultaneously detecting all updates to the source device, 11. This allows updated tracks to be migrated. Using the Innovation Data Processing, Inc. FDRPAS (FDR Plug and Swap) data migration engine, 13, the target device remains offline during the data migration.

Once data migration is completed, the Innovation Data Processing, Inc. FDRPAS (FDR Plug and Swap) data migration engine, 13, swaps the data storage devices, that is, it swaps the source and target data storage devices, 11 and 15.

## SIGNAL BEARING MEDIA AND ASSOCIATED PROGRAM PRODUCT

The invention may be implemented, for example, by having the data migration engine or tool, 13, under the control of one of more of the servers, 17 in FIG. 2, 17a or 17b in FIG. 3, or 17c or 17d in FIG. 4, execute a sequence of machine-readable instructions, which can also be referred to as code. These instructions may reside in various types of signal-bearing media. In this respect, one aspect of the present invention concerns a program product, comprising a signal-bearing medium or signal-bearing media tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform a method for migrating data from one storage medium to another storage medium by reading data from a source volume on a source data storage device, as a bit image of a logical volume; and writing the data to a target volume on a target data storage device as a bit image of a logical volume.

This signal-bearing medium may comprise, for example, memory in the servers, such as the illustrated 17 in FIG. 2, 17a or 17b in FIG. 3, or 17c or 17d in FIG. 4. The memory in the server may be non-volatile storage, a data disc, or even memory on a vendor server  
5 for downloading to one of the servers, 17 in FIG. 2, 17a or 17b in FIG. 3, or 17c or 17d in FIG. 4, for installation. Alternatively, the instructions may be embodied in a signal-bearing medium such as the optical data storage disc 30 shown in FIG. 5. The optical disc can be any type of signal bearing disc or disk, for example, a CD-ROM, CD-R, CD-RW, WORM, DVD-R, DVD+R, DVD-RW, or DVD+RW. Additionally, whether  
10 contained in a cluster with the server, 17, or elsewhere, the instructions may be stored on any of a variety of machine-readable data storage mediums or media, which may include, for example, a "hard drive", a RAID array, a RAMAC, a magnetic data storage diskette (such as a floppy disk), magnetic tape, digital optical tape, RAM, ROM, EPROM, EEPROM, flash memory, magneto-optical storage, paper punch cards, or any other  
15 suitable signal-bearing media including transmission media such as digital and/or analog communications links, which may be electrical, optical, and/or wireless. As an example, the machine-readable instructions may comprise software object code, compiled from a language such as "C++".

20 Additionally, the program code may, for example, be compressed, encrypted, or both, and may include executable files, script files and wizards for installation, as in Zip files and cab files. As used herein the term machine-readable instructions or code residing in or on signal-bearing media include all of the above means of delivery.

25 While our invention has been described with respect to certain preferred embodiments and exemplifications, it is not intended to limit the scope of the invention thereby, but solely by the claims appended hereto.